BrightStor[™]CA-Vtape[™] Virtual Tape System

Getting Started 2.0



BrightStor[™]CA-Vtape[™] Virtual Tape System

Getting Started





This documentation and related computer software program (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by Computer Associates International, Inc. ("CA") at any time.

This documentation may not be copied, transferred, reproduced, disclosed or duplicated, in whole or in part, without the prior written consent of CA. This documentation is proprietary information of CA and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of this documentation for their own internal use, provided that all CA copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the confidentiality provisions of the license for the software are permitted to have access to such copies.

This right to print copies is limited to the period during which the license for the product remains in full force and effect. Should the license terminate for any reason, it shall be the user's responsibility to return to CA the reproduced copies or to certify to CA that same have been destroyed.

To the extent permitted by applicable law, CA provides this documentation "as is" without warranty of any kind, including without limitation, any implied warranties of merchantability, fitness for a particular purpose or noninfringement. In no event will CA be liable to the end user or any third party for any loss or damage, direct or indirect, from the use of this documentation, including without limitation, lost profits, business interruption, goodwill, or lost data, even if CA is expressly advised of such loss or damage.

The use of any product referenced in this documentation and this documentation is governed by the end user's applicable license agreement.

The manufacturer of this documentation is Computer Associates International, Inc.

Provided with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227-7013(c)(1)(ii) or applicable successor provisions.

© 2001 Computer Associates International, Inc.

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

Chapter 1: Introduction

Introduction	1-1
Overview	1-2
How it Works	1-3
Why Software?	1-5
Efficient Tape Usage	1-6
Higher Quality, Lower Total Cost of Ownership (TCO)	1-6
High-Level Specifications	1-7
CA-Vtape Architecture	1-7
Virtual Volumes	1-7
Data spaces	1-8
Data Sets	1-8
Disk Buffer Pool	1-9
Virtual Allocation Management	1-10
New (Scratch) Allocations	1-10
Retrieving Existing Data Sets	1-11
Close (Tape Mark) Processing	1-12
Multi-Volume Output Processing	1-12
Multi-File Processing to Virtual Volumes	1-12
Data space Processing	1-13
Disk Buffer Pool Status	1-14
ICF Catalog Processing	1-14
Special Processing Conditions	1-15
Other Processing Conditions	1-15
Automated Virtual Volume Stacking	1-16
Stacking Process	1-16

DACD B (C. D. 1. 10) 1: 14	1-17
DASD Buffer Pool and Stacking Management	1-17
Disaster Recovery	1-18
Duplex Copies	1-18
Export Copies	1-18
Readability Outside CA-Vtape	
Tape Recycle	
Recycling Process	
Recycle Simulation and Reporting	1-20
What's New in Release 2.0	
Product Features	1-21
Documentation Changes	1-24
<u> </u>	
Chapter 2: Installation Considerations	
Installation Considerations	2-1
Hardware and Back-Store Requirements	
System/Install Requirements	
Chapter 3: Planning for Installation	
Volume Usage Summary Calculations	3-1
Volume Usage Summary Calculations	
Cache Size Alternatives to Preserve Cache Residence	3-2
Cache Size Alternatives to Preserve Cache Residence	3-2 3-2
Cache Size Alternatives to Preserve Cache Residence	3-2 3-2 3-3
Cache Size Alternatives to Preserve Cache Residence	
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types	
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze	
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data	3-2 3-3 3-3 3-4 3-4 3-4 3-5
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data JCL Setup	3-2 3-2 3-3 3-4 3-4 3-4 3-5
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data JCL Setup VMA and CA-Vtape Analysis and Reporting	3-2 3-2 3-3 3-4 3-4 3-5 3-5
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data JCL Setup VMA and CA-Vtape Analysis and Reporting JCL Setup	3-2 3-3 3-3 3-4 3-4 3-5 3-5 3-7
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data JCL Setup VMA and CA-Vtape Analysis and Reporting	3-2 3-2 3-3 3-4 3-4 3-5 3-5 3-7 3-7
Cache Size Alternatives to Preserve Cache Residence Tape Usage Projection and Estimated Cartridge Savings Understanding IBM's Volume Mount Analyzer Running CA-Vtape Modeling SMF Input Record Types How Much Data to Analyze VMA Consolidation of SMF Data JCL Setup VMA and CA-Vtape Analysis and Reporting JCL Setup Input Parameters	3-2 3-2 3-3 3-4 3-4 3-5 3-5 3-7 3-7 3-8

NOREPORTPHYVOL	3-9
SHOWINLINES	3-9
SHOWMIN(0)	3-9
SHOWMAX(99999)	3-10
CACHEFACTOR(0-100)	3-10
Example Output Reports	3-12
Heading and Input Control Statement Listings	3-12
Chapter 4: Customer Support	
Software that Manages eBusiness	4-1
Customer Support	
Getting Technical Support for CA-Vtape	

Index

Chapter

Introduction

The Getting Started introduces you to BrightStor CA-Vtape Virtual Tape System and provides instructions for installation.

To find out more about BrightStor CA-Vtape Virtual Tape System, see the *User Guide* and the *Messages Guide*.

Introduction

High capacity tapes. Tape stacking software. Tape mount management strategies. You've probably considered, or even tried these techniques to improve the efficiency of your tape storage environment. Now you're considering "virtual tape", bundled hardware and software systems intended to increase tape utilization and throughput.

Virtual tape is a great idea, but the hardware-based systems on the market today have several limitations. First, they are expensive and difficult to implement. Furthermore, they are often based on non-mainframe class components that severely limit their throughput and reliability.

You will have no guarantees that you can use new tape technologies when they arrive. And, most importantly, they require you to "bet your shop" on an expensive, proprietary system with limited, and costly, options for scalability.

CA-Vtape is a software-based virtual tape solution. It uses your existing tape and hardware to implement virtual tape with the mainframe-class reliability and performance you demand.

CA-Vtape is the most efficient, cost-effective virtual tape system available.

As your tape storage needs grow, CA-Vtape grows with you, without compromising your choice of hardware. In fact, you can upgrade resources to take advantage of capacity and performance advancements, regardless of the manufacturer.

CA-Vtape:

- Lets you use your existing mainframe tapes, drives, automated tape libraries, applications and tape management system
- Easily and cost-effectively scales to meet your growing storage needs without locking you into a single vendor
- Increases tape utilization for lower total cost of ownership
- Improves throughput
- Eliminates the duplicate costs ordinarily associated with using virtual tape in disaster recovery sites
- Is the least expensive virtual tape solution on the market

Overview

Software-based and hardware-independent, CA-Vtape lets you create a virtual tape environment from any combination of S/390 disk and tape hardware, including resources you already own.

You can also add or upgrade resources as needed to take advantage of advancements in capacity and performance, regardless of manufacturer.

How it Works

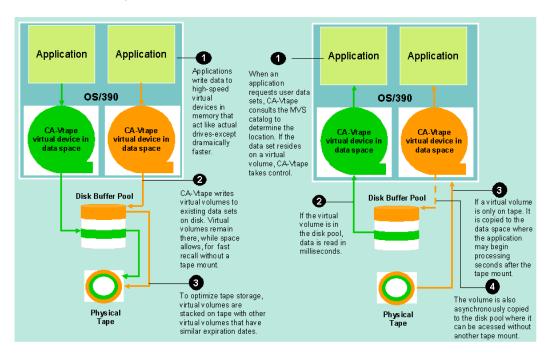
Applications write to high-speed buffers and data spaces, in OS/390 that emulate physical tape drives and volumes.

The application reads and writes to this data space at speeds equal to, or better than, native performance. From there, CA-Vtape takes over, writing the data to the disk buffer pool where it remains available for recall as long as space allows, see Figure 1-1 CA-Vtape in Action.

As soon as the virtual volume is completely written to the disk buffer, it is eligible to be externalized or stacked to physical tape. You have complete control of when externalization occurs. Stacking is based on predefined attributes such as data set names, DFSMS dataclass, and retention periods to group virtual volumes with similar expiration dates on the same physical tapes. This minimizes fragmentation, reduces wasted space on each cartridge, improves automation and ultimately makes it easier to manage vaulting and recovery. See Figure 1-2 DSN Filter List Display Panel.

Because recalls are typically made soon after a data set is written, having the virtual volume in the disk buffer is not only fast, it reduces tape mounts. However, when the virtual volume must be restored from physical tape, CA-Vtape transfers the first few blocks of the requested data to the data space. The application can access and begin processing a few seconds after the tape mount while the rest of the virtual volume is restored to disk asynchronously in the background. Subsequent access to that volume occurs without waiting for a physical tape mount. See Figure 1-1 CA-Vtape in Action.

FIGURE 1-1 CA-Vtape in Action



Why Software?

A software solution delivers a number of advantages. First, CA-Vtape works with any S/390 disk and tape hardware, with or without an automated tape library (ATL). So you are not locked into a hardware vendor or proprietary system. You can get the most from your existing resources and grow with technology that suits your needs.

Second, hardware systems are only available in a few preset capacities and price points. CA-Vtape lets you expand in increments that meet the tape-processing demands of your business.

Finally, you can easily evaluate CA-Vtape onsite, simply load a tape and go. Before you can truly evaluate a bundled system, you must commit time and resources to install the system. There are heavy crates, floor space considerations and environmental concerns, all before you can turn it on.

FIGURE 1-2 DSN Filter List Display Panel

```
Menu List Mode Functions Utilities Help
                CA-Vtape DSN Filter List Display Panel Row 1 to 9 of 22
Command ===> _
Dataset filter information. To Create one type ADD or enter PF11.
                                                        I/X Output Group
    Dataset
    ATTPE01.GRP01.*
    ATTPE01.GRP03.*
    PSS.A300.TOWLE02.*
   QAPROD.TSTGRP01.*
QAPROD.TSTGRP02.*
QAPROD.TSTGRP03.*
QAPROD.TSTGRP04.*
    RIDOS01.QAVTAPE.*
                                                        Т
                                                             1.1
   TANJO02.TSTGRP51.*
F1=Help F2=Spl:
F9=Swap F11=Add
              F2=Split F3=End F5=Refresh F7=Up
                                                                   F8 = Down
```

Efficient Tape Usage

CA-Vtape maximizes the effectiveness of tape and improves ATL slot management.

By automatically stacking virtual volumes on physical tapes, CA-Vtape virtually eliminates abends and extended recall delays associated with conventional stacking products.

The result? More data is contained on fewer tapes — as much as a 300 percent improvement over non-stacked tapes so more cells and devices are available to meet the demands of your growing operation.

Another way that CA-Vtape improves tape usage is through its tape recycle capability. This utility consolidates fragmented physical tapes into a few fully-used tapes.

Higher Quality, Lower Total Cost of Ownership (TCO)

Many hardware-based virtual tape systems are built from components that fall short of mainframe performance standards. CA-Vtape delivers mainframe-class reliability and higher tape utilization at a low cost. That's because CA-Vtape uses DASD from your existing system for its intermediate disk buffer pool.

Through efficient tape usage, CA-Vtape will actually help you eliminate tape devices, including ATLs, and reclaim floor space. The result is a very low TCO for CA-Vtape, substantially lower than any hardware-based virtual tape system and with the quality and performance you expect from your mainframe.

High-Level Specifications

Emulation: 3480 or 3490

Maximum number of virtual devices: 200

Maximum number of virtual volumes: 500,000*

Maximum size of disk buffer pool: 50 TB*

CA-Vtape Architecture

The CA-Vtape virtual tape environment is composed of virtual devices, virtual volumes, data spaces, data sets, and the disk buffer pool. Through the main menu and ISPF interface, you can drill down to manage CA-Vtape objects.

CA-Vtape creates virtual tape devices inside OS/390. Because CA-Vtape emulates the microcode that controls 3480 and 3490 tapes, OS/390 acts as if the device actually exists. Applications read from and write to these devices exactly as they would interface with a real unit.

Virtual Volumes

Any tape management system will manage virtual volumes exactly like a physical tape. You define the volume range to the tape management system and to CA-Vtape during installation. CA-Vtape will synchronize scratch volumes with the tape management system as well as verify the volume status and ownership. Virtual volumes may reside either in the disk buffer pool, on a physical tape stacked with other virtual volumes, or in both.

^{*} For each CA-Vtape complex.

Data spaces

OS/390 data spaces act as high-speed memory buffers for CA-Vtape. When a mount occurs, CA-Vtape creates a dedicated data space for each virtual device. As applications read or write a virtual volume, CA-Vtape moves data between the data space and the disk buffer pool.

Data Sets

CA-Vtape data sets are VSAM linear data sets (LDS) that contain all the data for CA-Vtape virtual volumes including control and indexing information. These LDSs are preallocated and cataloged in the disk buffer using standard VSAM services.

Once the virtual volume is copied to physical tape, the virtual volume LDS may be reclaimed. The reclaim is triggered by buffer pool management algorithms. It reclaims space from the least-recently used (LRU) virtual volume first.

Disk Buffer Pool

CA-Vtape uses a disk buffer pool in DASD to contain data before and after it is written to physical tape. As long as a virtual volume remains in the buffer pool, the volume will be recalled from there. You can set the buffer pool size and location during installation. After defining the initial buffer pool size, you can dynamically extend it across any combination of storage devices in your environment. When the disk buffer pool reaches the maximum size, CA-Vtape reclaims space by deleting the leastrecently used virtual volume LDS that has already been copied to a physical tape.

Large tape data sets are segmented into multiple virtual volumes. This allows those that have completed processing to be copied or destaged to physical tape while subsequent volumes for that data set continue processing. As a result, very large tape data sets will not completely consume the disk data set pool or constrain other tape data sets.

FIGURE 1-3 CA-Vtape Main Menu

```
CA-Vtape Main Menu
                                                                                _____
Option ===>
     Select one of the following:
     1 Group List Display output group information
2 Tape Devices Display SVTS related tape devices
3 Virtual Volumes Manage virtual volumes
4 Dataset List Manage the dataset filter list
5 Dataclass List Manage the Dataclass filt.
     Enter PF3 to exit
PF 1=HELP
                           2 = SPLIT
                                                 3 = END
                                                                       4 = RETURN
                                                                                            5 = RFIND
PF 6 = RCHANGE
                           7 = UP
                                                 8 = DOWN
                                                                       9 = SWAP
                                                                                           10 = LEFT
```

Virtual Allocation Management

CA-Vtape lets your applications write to virtual tape devices just as they would to physical devices. Data is passed from the user's application to virtual devices in OS/390 data spaces. These data spaces pass the data to data sets in the disk buffer pool that are exclusively allocated to the virtual volumes.

New (Scratch) Allocations

CA-Vtape is invoked when the system selects the appropriate device for the application. The data set name or DFSMS dataclass is compared to CA-Vtape's filter list control tables. If it is included, CA-Vtape is invoked.

For scratch tape mounts, CA-Vtape selects the virtual volume serial number from an internal list of scratch candidates which are synchronized by the tape management system. To satisfy the virtual volume scratch allocation, CA-Vtape selects 10 VSAM LDSs from the available cell list and opens the virtual volume. To ensure the most economical use of space, each LDS is 1/10th of the virtual volume size. This also enables each virtual volume, and the buffer pool, to be distributed across multiple storage devices for enhanced performance.

After you create a new virtual volume, its information appears, along with all the virtual volumes, on the CA-Vtape display panel. There you can see which virtual volumes are in the buffer pool, their size, data set names, virtual VOLSERs, and status.

FIGURE 1-4 Virtual Volume Display Panel

<u>F</u> ile <u>H</u> elp				
Command ===>	CA-Vtape Virtual Volume Display Panel Ro	ow 91 to	100 o	f 10
Cache Limit(MB) 8040 Cache Inuse(MB) 6720 MB	Freeabl	e 6720	
Use S to dis	play additional virtual volume data			
Volid	Tape DSN	# MB	Allo	RES
	RIDOSO1.QAVTAPE.G9	236		Y
_ 100491	RIDOSO1.QAVTAPE.K9	200		Y
_ 100492	RIDOSO1.QAVTAPE.E9	236	240	Y
_ 100493	RIDOSO1.QAVTAPE.H9	236	240	Y
_ 100494	RIDOSO1.QAVTAPE.I9	236	240	Y
100495	RIDOSO1.QAVTAPE.F9	12	40	Y
_ 100496	ZOBKU01.SVTGRP12.SORTOUT	16	40	Y
100497	ZOBKU01.SVTGRP12.SORTOUT	16	40	Y
_ 100498	RIDOSO1.QAVTAPE.D9	236	240	Y
100499	RIDOSO1.QAVTAPE.C9	236	240	Y
	F2=Split F3=End F5=Refresh F7=Up	TO -	D	

Retrieving Existing Data Sets

When an application requests an existing data set, OS/390 finds it in the JCL or ICF catalog and provides the requested virtual volume to CA-Vtape.

If the data set is in the buffer pool, CA-Vtape provides it to the application immediately. If not, it must be recalled from physical tape. Once selected, OS/390 positions the tape in the correct file sequence, mounts the virtual volume and begins data transfer. After the first blocks of data are copied to the data space managed by CA-Vtape, the requesting application may access and process data immediately. The rest of the data set is staged to the disk buffer pool at high speed, where it's available for future access.

Meanwhile, the user continues concurrent processing from the data space.

Retrieving different virtual volumes that are stacked on the same physical tape won't generate ENQ contention or abends. CA-Vtape will release the tape only after all the outstanding requests are satisfied. Even if there are multiple recalls from the same tape, only one mount is required.

If the application performs a "block space forward" search or "locate a block" command during the reading process, CA-Vtape needs only milliseconds to search the index in memory, then access the data in storage or from the buffer pool. CA-Vtape also provides advantages during readbackward processing. This process is performed at processor storage or DASD speeds without the limitations of autoblocking and other mechanical processes.

Close (Tape Mark) Processing

All data is synchronized and written to the virtual volume in the disk buffer pool at tape mark process time. As a result, all of the data is on disk before the virtual volume is dismounted.

This ensures data integrity between the application and the virtual volume disk data set.

Multi-Volume Output Processing

CA-Vtape forces an "End Of Volume" condition during output processing when your predefined maximum file size is met. In small buffer pools, small volume sizes will force multi-volume segmentation early so the initial virtual volumes can be moved to physical tape while the job continues on subsequent virtual volumes.

Multi-File Processing to Virtual Volumes

CA-Vtape will automatically stack virtual volumes to physical tape. If multi-file processing is required, CA-Vtape supports up to 255 files per virtual volume. If an attempt is made to create more than 255 files per virtual volume, CA-Vtape will enforce a volume switch to ensure optimum recall and recovery performance.

Data space Processing

The data being written to a virtual device is saved in VSAM LDSs in the disk buffer pool that are the save areas for the data space. This allows data to be written to the virtual volume asynchronously with data being written to the disk buffer. The application does not wait for each block of data to be written; however, all data must be written to the disk buffer before processing completes.

When multiple virtual devices process concurrently, multiple data space tasks concurrently write the data to the buffer pool to avoid "single thread" processing. Since the data spaces may accept data from applications faster than their buffer pool can accept it, CA-Vtape detects delays and controls the virtual device by presenting "busy conditions." This prevents paging of virtual storage and impacting other applications.

Disk Buffer Pool Status

Using a console command or through a CA-Vtape ISPF screen, you can easily determine the current buffer pool status, which jobs are using virtual devices, the status of each job, which physical tapes are being used for stacking, and the percentage of work completed for those tapes.

FIGURE 1-5 Tape Device Status Display Panel

```
File Help

CA-Vtape Tape Device Status Display Panel Row 1 to 9

Command ===>

Cache Limit(MB) 8040 Cache Inuse(MB) 7920 MB Freeable 216

The following device list contains information on virtual tape and physical devices allocated to CA-Vtape.

Volume Dev# Status Devt Johname # SIOS SEQ# Phy/Virt % 507499 0748 I / 0 349S SVTSXE61 22 6 Physical 0 100511 3500 I / 0 349S RIDOSG14 1679 1 Virtual 0 100513 3501 I / 0 349S RIDOSG14 1679 1 Virtual 0 104 349S 3501 I / 0 349S RIDOSG14 1894 1 Virtual 0 104 3503 349S 3 0 Virtual 0 10516 3505 I / 0 349S RIDOSH14 894 1 Virtual 0 104 3503 349S 3 0 Virtual 0 100516 3505 I / 0 349S RIDOSH14 590 1 Virtual 0 100517 3506 I / 0 349S RIDOSH14 590 1 Virtual 0 100517 3506 I / 0 349S RIDOSH14 590 1 Virtual 0 100517 3506 I / 0 349S RIDOSH14 319 1 Virtual 0 100517 3506 I / 0 349S RIDOSH14 319 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 319 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 319 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 319 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 3179 1 Virtual 0 100510 3507 I / 0 349S RIDOSH14 575 REFERENCE FF / 0 F / 0 100510 A / 0 1005
```

ICF Catalog Processing

Requests for a virtual volume are passed to CA-Vtape for processing. Applications continue to reference the original data set name without modification of the application or JCL. If the virtual volume is still in the DASD buffer pool, CA-Vtape will use its internal catalog, the Global VCAT, to select the appropriate virtual volume VSAM LDS data sets. If the virtual volume is not in the buffer pool, the ICF Catalog is searched for the appropriate data set name, that is, searching for either the primary or duplex tape.

Virtual volumes residing on physical tapes take on the following naming structure; HLQ.VVE.Vxxx.PRIMARY and/or HLQ.VVE.Vxxx.DUPLEX, where xxx is the virtual volume VOLSER.

Special Processing Conditions

- Uncataloged data sets on the virtual volumes are supported. VOL=SER references the requested virtual volume instead of a catalog lookup. Other processing remains unaffected.
- Unit Affinity Processing is supported when all of the data sets referenced by the Unit Affinity are in the CA-Vtape control tables.
- Reference Backward Processing (Refer Backs) is supported if all of the data set names invoke CA-Vtape.
- Data set MOD Processing is supported as DISP=MOD data sets. The virtual volume will be recalled and extended as desired.
- Concatenation Processing may refer to both data sets serviced by CA-Vtape and data sets not serviced by CA-Vtape in the same step.

Other Processing Conditions

CA-Vtape will honor JCL parameters including OPTCD=W, OPTCD=C, and OPTCD=Z.

Automated Virtual Volume Stacking

With CA-Vtape's automated virtual volume stacking, you can organize related data for the best access and tape utilization.

FIGURE 1-6 Group Display Panel

							Displ				 Row 1	
omn	and ==	= >		_ H - V C C	ape o	roup	DISPI	ay ro	neı		KOW I	
his	s panel	displays	the	e def:	ined	Group	s.					
ax	Drives	. : 2	D:	rives	In U	se .	: 0					
	Group	Id WRK Q	' D	Prim	Eso	Dup	Eso	Ехр	Eso	Short	$M \in d$	Loi
_	1	N		SILO						7	21	> 0 (
_	2	N		SILO		SILO)			7	21	> 0 (
_	3	N		SILO				SILO		7	21	> 0 (
	4	N		SILO		SILO)	SILO		7	21	> 0 (
_	1 1	N		SILO						7	21	> 0
	12	Y		SILO		SILO)			7	21	> 0 (
	13	N		SILO				SILO		7	21	> 0
	1 4	N		SILO		SILO)	SILO		7	21	> 0 (
F1=	Help	F2=Sp1	it	F3:	End		F5 = Ret	fresh	F 7	= U p	F8 = 1	Down
F 9 =	:Swap	F11 = Add										

Stacking Process

The CA-Vtape stacking format allows simultaneous access to multiple virtual volumes on a physical tape without abends or extended delays.

When virtual volumes are created or updated, the virtual volume resides on disk but must still be copied to physical tape. Once the virtual volume is dismounted, it is scheduled to be copied from the buffer pool to physical tape. To improve physical tape utilization, virtual volumes are stacked together based on their stacking group and retention period.

Stacking Groups

Stacking groups allow data sets with similar characteristics and requirements to be stacked together. Through your predefined filters, CA-Vtape automatically controls the group assignment and stacking process with minimal ongoing management. Data sets are stacked by:

- Expiration Date Stacking Groups have three subgroups; short, medium and long, that automatically separate the data sets into different physical tapes based on expiration date. You may define different retention period lengths for each subgroup. This minimizes fragmentation of freespace on cartridges.
- Desired Esoteric, that is, automatic vs. manual device.
- Desired Storage Location, that is, local vs. disaster vault.
- A user-assigned separator group, that is, a user may wish to ensure the second copy of a database log is on a different physical tape than the primary copy.

DASD Buffer Pool and Stacking Management

CA-Vtape automatically monitors the space in the buffer pool. When your specified threshold is reached, an alert is sent to the system console. There, you can dynamically change the threshold value, assign additional storage to the buffer pool, or stack the data. You may also select the number of physical tape drives dedicated for output.

You can independently hold or release the stacking process using any combination of Groups or Subgroups. You have total control over when and how this occurs via any scheduler or automation tool. Or you can enter commands from the system console.

Disaster Recovery

Access to data is a must, whether for disaster recovery or readability outside your tape system. Here's how CA-Vtape helps.

With CA-Vtape, disaster recovery is simple. You don't need duplicate or specialized virtual tape hardware at a remote disaster recovery location. You can quickly activate CA-Vtape and begin using and accessing the data it manages.

In addition to creating a PRIMARY copy, CA-Vtape allows several options for output, including DUPLEX and EXPORT, that you can vault for recovery.

Duplex Copies

Duplexing creates a second, stacked physical tape copy of the data sets in a local or remote site for media or disaster recovery. As with the primary tape data set, the duplex copy is stacked in the CA-Vtape format for optimum tape utilization. See the section "Export Copies" in this chapter for a non-proprietary option. Duplex copies may be used in any OS/390 system containing CA-Vtape, for example, at an offsite location. And these copies can even be on different media than the primary copy. This way, duplexing can provide a vault copy for disaster backup and recovery processing and for local recovery as well.

Export Copies

Occasionally, physical tapes must be directly readable without CA-Vtape, for example, for readability by an outside data center. To satisfy this need, you can create an *Export* copy. This separate, uncataloged copy of a virtual volume is copied in native MVS format, so CA-Vtape is not required to access the data on this tape.

Readability Outside CA-Vtape

CA-Vtape includes a utility that reads primary or duplex tapes stacked by CA-Vtape and extracts the data in a user-readable format. This utility may be retained without a license to ensure data accessibility with or without CA-Vtape.

Tape Recycle

Tape recycling is critical to ensure your environment benefits from all that virtual tape has to offer.

Recycling Process

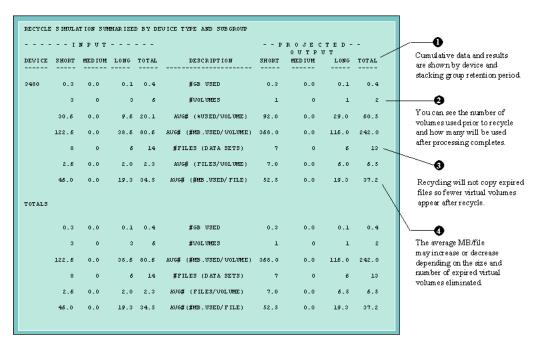
To ensure ongoing efficiency, the CA-Vtape Recycle function consolidates fragmented physical tapes to recover free space. Recycled tapes can be returned to the tape management system scratch pool and reused. You can even simulate the proposed recycle to forecast how many tapes are involved and the potential savings before you commit time to the process.

Recycle is a tape-to-tape batch utility that supports concurrent executions and provides comprehensive filtering for a streamlined process. Additionally, the Recall Detection facility ensures production processes are not negatively impacted or delayed. If a recall is detected, CA-Vtape provides the application with the physical tape containing the requested virtual volume. CA-Vtape will continue the recycle process with another physical tape.

Recycle Simulation and Reporting

To help determine the results of recycle or simulated recycle, CA-Vtape reports on the activity. You can see how many physical tapes are affected and the outcome. Note that the recycle report refers to files, meaning virtual volumes, and volumes, meaning physical tapes.

FIGURE 1-7 Recycle Report



What's New in Release 2.0

This summary includes Release 2.0 changes to CA-Vtape and to its documentation.

Product Features

CA-Vtape Release 2.0 is a major release for the product and includes new functionality in several areas as requested by clients. This release demonstrates Computer Associates continued commitment with our clients and the intention to remain an important player in the virtual tape market.

The following features are new and available in Release 2.0:

Dynamically changing group definitions

Group definitions reside in PARMLIB and a refresh command is provided. Each individual group attribute can be dynamically changed and refreshed.

Add UCB ranges during installation

UCB definitions reside in PARMLIB and UCB ranges like 0F00:0F0F will be allowed. Using the new parameters to define UCBs, clients can also tell CA-Vtape to automatically bring the virtual drives online during startup or to keep them offline as is currently done in CA-Vtape Release 1.2.

Parameter library

PARMLIB is implemented and some of the actual definable parameters and thresholds are supported in this release. Future releases will add other parameters. The goal is to have all the parameters, definitions and customization as members of the PARMLIB providing the flexibility of editing, replacing and in some cases dynamically refreshing them. For more details about this new feature see the appendix "Using the PARMLIB Feature" in the *User Guide*.

Only one SVTSIO address space

This release is using a single SVTSIO address space. No changes are needed in the installation or implementation.

Enhanced externalization and recall

Recalls and externalization in Release 2.0 are running in a separate address space with individual controls allowing proper management of the physical tapes and also independence from the SVTS address space. With this feature recalls will be asynchronous with the requestor and can continue running if the original application requiring the virtual volume is no longer in the system. New console commands are provided to display back end activity, start/stop recalls, stop a particular task or even to refresh the entire back end support without effecting the front end. Enhanced support for multi-system recall detection is also included. For more details about this new feature see the appendix "Enhanced Externalization and Recall Feature" in the *User Guide*.

Compression at data space level

Hardware compression algorithms provided by OS/390 can be used to compress in the dataspace the data transferred from the applications. A user-definable threshold is now implemented to define the compression level desired for the CA-Vtape system or each individual group. The data will flow in compressed format from the dataspace to the DASD buffer and to the physical tapes. This functionality will help to reduce I/O load, DASD buffer space, and improve performance and throughput. Data will be uncompressed only when the application receives the data back during a recall. Compressed and uncompressed virtual volumes can be combined on the same physical tape. Compression can be activated globally or at group level.

Improved data set filter list pattern mask

Data set filters and Dataclass lists reside in PARMLIB and a complete set of pattern masks were added for the data set filters using CA-Vantage standards. CA-Vantage, CA's Storage Resource Manager, features the Tape Resource Manager component that also includes an easy to use, Windows based CA-Vtape interface.

Use of BLOCKID for physical tapes positioning

Used during Recalls, Externalization and Recycle, this new feature dramatically improves the elapsed time a tape unit uses positioning a physical tape.

Command to mark a physical tape "full"

Required for tapes being sent offsite, this new command instructs CA-Vtape to not reuse the physical tape when externalization is restarted.

Standardize the PF Key Settings used in the ISPF panels

The PF key settings for some of the panels were changed to make them compliant with ISPF standards.

"Recall only" virtual scratch ranges

Users are now able to dynamically add virtual volume ranges with a "use only for recall" attribute. This allows two or more active but not interconnected CA-Vtape sites to act as mutual Disaster Recovery sites. Each one must define the other's virtual volume scratch ranges as "recall only" and then they will be able to exchange and recall any physical tape in any site at any time without the need to also recover VCAT information.

Enhanced Global Vcat Recovery

New options were added to the GVCAT from the BSDS recovery process allowing a faster and easier recovery in the same or a different site, even when part of the CA-Vtape DASD buffer is no longer available.

64-bit support

z/OS and the new z900 IBM mainframes running in 64-bit are supported by Release 2.0.

Documentation Changes

The documentation set has been engineered to take advantage of the latest technology for online viewing, keyword searching, book marking, and printing. This set contains a hardcopy of the *Getting Started*, the set also contains the *Getting Started*, the *User Guide* and the *Messages Guide* in Adobe Acrobat Reader format on the CD.

- The *Getting Started* guide is new and provides the user with an overview of the system, what's new in Release 2.0, system requirements and gets you started with installation of the product.
- The *User Guide* was split up into 2 manuals;
 - the User Guide

This guide contains detailed installation steps and how to use and troubleshoot CA-Vtape and its components.

the Messages Guide

This guide is an alphanumeric listing of the messages produced by CA-Vtape and its components. This guide advises you what prompted CA-Vtape to give a message and the action you can take.

Chapter

Installation Considerations

Installation Considerations

Hardware and Back-Store Requirements

- Any S/390 compatible disk device (mainframe-class RAID devices recommended)
- Any S/390 compatible tape or automated tape library (ATL is not required)

System/Install Requirements

- OS/390 2.6 and above and z/OS 1.1 and above running in 31-bit or 64-bit mode
- HCD definition of the virtual tape drives (software only)
- TSO/ISPF access
- Access to SYS1.PROCLIB or similar library in the JES2 proclib concatenation
- TopSecret or similar security features and permission to grant authority to the CA-Vtape software
- Any tape management system

Chapter

Planning for Installation

Before you install CA-Vtape, you can evaluate your system to see the storage and other improvements you can expect. The CA-Vtape Modeling Utility uses input from IBM's Volume Mount Analyzer (VMA) reports to create CA-Vtape specific installation analysis reports and implementation guidelines. Using the VMA ensures your analysis of an installation's tape usage is consistent with other tape analysis projects your installation may have already done. In addition to normal VMA utility output, the CA-Vtape modeler produces the following reports:

Volume Usage Summary Calculations

This report shows a breakdown of the number of tape volumes in your installation and the number of volumes in megabyte size ranges. This information gives you information about how CA-Vtape can be effective in consolidating tape volumes and statistics on what sizes of data sets you want CA-Vtape to manage.

Cache Size Alternatives to Preserve Cache Residence

This report shows projections of the size of DASD cache to be set up for different size ranges of CA-Vtape volumes. It helps you to understand the amount of DASD cache you need depending on how long you want virtual volumes to remain on quickly accessible cache storage.

Tape Usage Projection and Estimated Cartridge Savings

This report projects the potential tape cartridge savings you can achieve when CA-Vtape is used to manage different size ranges of data sets. It demonstrates the dramatic savings obtained by implementing CA-Vtape.

Understanding IBM's Volume Mount Analyzer

If you are unfamiliar with IBM's tape VMA, read the documentation for the utility and make yourself generally aware of setup and options. CA-Vtape analysis reports only use a small part of the overall capabilities of VMA and we are not attempting to reproduce the operational documentation here. Refer to the IBM publication, *DFSMS/MVS Using the Volume Mount Analyzer* – Document SC26-4925-01 for more information.

VMA and CA-Vtape analysis reports are run in a two job process.

- The first job uses the VMA extract to filter and extract large amounts of SMF data to create a consolidated file to be used for quick repeated analysis of the extracted data. Since your SMF log data can be very large, the separate processing of the data for consolidation ensures you only need to scan your SMF data once even though you want to run several variations of analysis reports.
- The second job of the analysis is to input the consolidated extract file and create a VMA output file to be processed by the CA-Vtape Modeler Utility. The second step in this job runs the CA-Vtape analysis and creates the three modeler analysis reports.

Both steps of VMA allow input parameters to be specified for filtering your SMF data. You can include or exclude data by jobname, data set name, date, time, unit and other parameters. Refer to the VMA documentation for specific information about filtering data.

Running CA-Vtape Modeling

You must first determine the location of your SMF log data sets and find the volumes containing records for the period of time your analysis is to be run. You can run the analysis for any period of time and for any number of input LPARs. Include information that reflects all the data you are considering for CA-Vtape implementation.

SMF Input Record Types

IBM's VMA requires a minimum level of SMF record types as part of your input log data. OS/390 SMF related system parameters determine what SMF data you are gathering on your system. If you are not sure of your SMF data contents, examine the SMF setup specifications or contact the person responsible for setting SMF specifications for your installation.

Following are the SMF record types required. For more information, refer to the IBM VMA documentation.

04, 05, 14, 15, 21, 30 - sub-type (4, 5, 34, 35), 34, 35

How Much Data to Analyze

Comparing the following subsets of data helps you better understand how the overall average hourly tape usage differs from peak average tape usage.

- Consider how much SMF data analysis is necessary to get a good picture of the makeup of your tape storage inventory.
- Consider at least a full 30 days data to make sure the analysis reflects monthly peak processing cycles.
- Run several sets of data to differentiate the characteristics of your system during peak processing times versus non-peak times.

VMA Consolidation of SMF Data

The first step in the analysis consolidates the raw input SMF data. Figure A-1 shows a scaled down sample JCL for running this consolidation. This JCL is member VMAJCL1 in the CA-Vtape INSTALL library. A more complete version of this JCL is provided by IBM in SYS1.SAMPLIB as member GFTAXTRP.

JCL Setup

Copy the JCL member VMAJCL1. Modify the JCL considering the following items:

- Add a job card.
- Specify the HLQ parameter with a data set name high level qualifier used to prefix the generated output data set name.
- 3. Check all marked statements for proper space parameter information. The VMA manual makes specific recommendations regarding these size specifications. Depending on how much data you process, this job can run for an hour or more so it is important to provide enough space to avoid abends because of lack of space.
- Specify the dsname and volser information for all input SMF data you want to process.
- 5. Check the output data set specification to ensure it is valid and is on a volume where it can be saved for later processing.
- 6. Due to long processing times, confirm the TIME parameter on the JOB and STEP statement to ensure your job is not automatically terminated for lack of time.
- 7. You can specify VMA statements to limit collection of data to specific time periods. For example, for all the days in one month. Refer to the VMA documentation for details. Input parameters are not required so the XTRCNTL DD statement may refer to an empty list of input.

After all specifications are made, submit the job for processing.

FIGURE 3-1 Contents of JCL member VMAJCL1

```
//jobname JOB (....), 'USER'.....
//*-----
//* VMA STEP1 of 2 SMF EXTRACT - CREATE CONSOLOIDATED INPUT FOR STEP2
//* Refer to IBM Publication below for instructions for this JCL:
//* DFSMS/MVS Using the Volume Mount Analyzer - Document SC26-4925
//* USES SMF RECORDS 14, 15, 21, AND 30 (SUBTYPE 4 AND 5).
//* ALSO USES 4, 5, 34, 35
//* IBM PROVIDES SAMPLE JCL AS MEMBER GFTAXTRP FROM SYS1.SAMPLIB
//* Change <-- marked statements below as appropriate for your system
//*----
//GFTAXTR PROC HLQ=, * GFTAXTR FILE HIGH LEVEL QUALIFIER
                               RUN=, * GFTAXTR FILE LOW LEVEL QUALIFIER
//
                               REGXTR=6144K, * DEFAULT REGION FOR 'GFTAXTR'
//
                               SMFIN=, * INPUT SMF dataset NAME
//
                               UNIT=SYSDA * DEFAULT DASD ESOTERIC UNIT NAME
//
//XTRACT EXEC PGM=GFTAXTR, REGION=&REGXTR, TIME=160
//SMFTN DD DTSP=SHR, DSN=&SMFIN
//SMFIN
                     DD DISP=SHR,DSN=&SMFIN
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SORTDIAG DD SYSOUT=*
//SORTDIAG DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//XTRCIN DD UNIT=&UNIT, DSN=&HLQ..GFTAXTR.&RUN,
                    DISP=(NEW,CATLG,DELETE),
//
                        RECFM=VB, SPACE=(CYL, (40,30), RLSE) <-- OUT FILE SPACE?
//
//XTRCOUT
                   DD UNIT=&UNIT, DSN=&&XTRCOU,
                        RECFM=VB, SPACE=(CYL, (40, 30), RLSE) <-- WORK SPACE?
//XTRCWK01 DD UNIT=&UNIT, SPACE=(CYL, (20, 30), ,CONTIG, ROUND) <-- SPACE?
//XTRCWK02 DD UNIT=&UNIT, SPACE=(CYL, (20, 30), ,CONTIG, ROUND) <-- SPACE?
//XTRCWK03 DD UNIT=&UNIT,SPACE=(CYL,(20,30),,CONTIG,ROUND) <-- SPACE?
    PEND
//
       EXEC GFTAXTR, HLQ=your.hlq, RUN=XTRACT#1
                                                        <-- OUT HLQ?
//* MODIFY THE BELOW //XTRACT.SMFIN DD TO REFER TO SMF LOG FILES
//XTRACT.SMFIN DD DISP=OLD,
             DSN=your.installation.smf.data.files,UNIT=xxxx,
//
//
                        VOL=SER=(...., SMF INPUT?
//XTRACT.XTRCNTL DD * No input statements required
```

VMA and CA-Vtape Analysis and Reporting

The second job in the analysis processes the consolidated input file created by job 1 and creates the three modeler reports. Scaled down sample JCL for running this consolidation is provided. This JCL is a member of VMAJCL2 in the CA-Vtape INSTALL library. A more complete version of this JCL is provided by IBM in SYS1.SAMPLIB as member GFTAVMAP.

JCL Setup

Copy the JCL member VMAJCL2. Modify the JCL considering the following items:

- Add a job card.
- Modify the &TAPES substitution parameter to indicate the estimated tape mounts you believe your SMF input data contains. The JCL is set up to use this value for determining work file space allocations. The second tape value is used to determine the secondary extent allocation size and should reflect 10-50% of the value of the first parameter.
- 3. The SYSEXEC DD statement in step 2 refers to the PDS library containing the modeling program. This library should contain the VTMODEL member which is the CA-Vtape modeling program. If you received this package from other than the CA-Vtape distribution tape, you may need to create or select a library to use and place the VTMODEL program in that library.
- 4. You may want to specify VMA statements to limit processing subsets of your tape data. For example, to include or exclude tape volumes by jobname, unitname, data set name or dates. Refer to the VMA documentation for details. The sample is provided with 2 input statements. The first is a sample to exclude all data for CA-Vtape virtual volumes. The second required statement indicates the VMA analysis output to be generated for CA-Vtape modeling.

5. You can provide CA-Vtape analysis control statements as explained in the section Example Output Reports in this chapter.

After all changes are made, submit the JCL for processing.

Input Parameters

The SYSIN DD for the STEP2 EXEC IRXJCL step provides optional parameters to control analysis output and provide parameters to customize analysis to your specifications.

Any number of parameter statements can be included in the SYSIN stream.

Parameters can be specified on separate statements, or can be separated by spaces and combined on the same input statement. Any statement with an asterisk (*) in column 1 is considered a comment statement and is ignored. Blank statements are also ignored.

The following parameter can be included:

VSIZE(400) or VSIZE(800)

You can specify that you use 400mb or 800mb virtual volume sizes for CA-Vtape installation. The value is used to calculate cache sizes and to determine the number of virtual volumes your data requires.

OUTPHYVOLSIZE (800)

Output Physical Volume Size

Indicate the amount of uncompressed megabytes you expect your output physical volumes to hold. This is used to estimate the number of physical backstor volumes that are needed to hold input virtual volumes. 800 indicates 800MB or .8 GB. If you compress output data on tape, specify a size that is larger than the physical capacity of the output tape cartridge size.

NOREPORTUSAGE

The Megabyte Range Usage Report #1 is always created from the analysis unless this parameter is used to eliminate it from the output.

NOREPORTCACHE

The Cache Report #2 is always created from the analysis unless this parameter is used to eliminate it from the output.

NOREPORTPHYVOL

The Physical Volume Report #3 is always created from the analysis unless this parameter is used to eliminate it from the output.

SHOWINLINES

Use to show the VMA input lines on the output utility report file. This parameter is primarily used for diagnostic purposes. When reporting modeling utility questions back to technical support, include the SHOWINLINES parameter to provide all possible documentation necessary for resolving your situation.

SHOWMIN(0)

Each report shows categories of output ranges by megabytes. It shows how managing different sized data impacts CA-Vtape implementation. You can eliminate small report size ranges by specifying a Minimum value to be displayed.

SHOWMIN(200) shows only size ranges above 199 megabytes.

SHOWMAX(99999)

Each report shows categories of output ranges by megabytes. It shows how managing different sized data impacts CA-Vtape implementation. You can eliminate large report size ranges by specifying a Maximum value to be displayed.

SHOWMAX(1000) shows only size ranges below 1001 megabytes.

CACHEFACTOR(0-100)

The percentage value specified by this parameter adds a percentage factor to the cache estimates shown on the Cache report. If a CACHEFACTOR(20) is specified then the values shown reflect an additional 20% of estimated space. This factor can be specified to account for cache space that is needed to hold tape recall requests for data sets that were written prior to the start date the VMA input SMF data was recorded. The default value for CACHEFACTOR is 20.

FIGURE 3-2 Contents of JCL member VMAJCL2

```
//jobname2 JOB (....), 'USER'.....
//*-----
//* VMA JOB 2 of 2 VMA AND CA-VTAPE ANALYSIS
//*-----
//* Refer to IBM Publication below for instructions for this JCL:
//* DFSMS/MVS Using the Volume Mount Analyzer - Document SC26-4925
//* IBM PROVIDES SAMPLE JCL AS MEMBER GFTAVMAP FROM SYS1.SAMPLIB
//* Change <-- marked statements below as appropriate for your system
//*-----
//GFTAVMA PROC HLQ=,
                                 * GFTAXTR FILE HIGH LEVEL OUALIFIER
              RUN=, * GFTAXTR FILE HIGH LEVEL QUALIFIER
REGVMA=6144K, * DEFAULT REGION FOR 'GFTAVMA'
//
//
//
               TAPES='20000,3000', * DEFAULT NUMBER OF TAPE MOUNTS
               UNIT=SYSDA, * DEFAULT ESOTERIC UNIT NAME
FILTERS=1000 * MAXIMUM NUMBER OF FILTERS
//
//VMA EXEC PGM=GFTAVMA, REGION=&REGVMA, PARM='FILT#(&FILTERS)'
//SORTIN DD UNIT=&UNIT,
               SPACE=(1000,(&TAPES)),AVGREC=U
//
//SORTOUT DD UNIT=&UNIT,
             SPACE=(1000, (&TAPES)), AVGREC=U
//SORTWK01 DD UNIT=&UNIT,
               SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
//SORTWK02 DD UNIT=&UNIT,
                SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
```

```
//SORTWK03 DD UNIT=&UNIT,
                 SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
//
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//VMACHART DD SYSOUT=*
//VMAINCL DD DUMMY, DCB=(DSORG=PS, RECFM=VB)
//VMAEXCL DD DUMMY, DCB=(DSORG=PS, RECFM=VB)
//VMACNTL
           DD DUMMY
//VMAFLTRS DD UNIT=&UNIT,
                 SPACE=(10, (&TAPES), RLSE), AVGREC=U
//XTRCIN DD UNIT=&UNIT,
                 SPACE=(1000.(&TAPES)).AVGREC=U
//XTRCOUT DD UNIT=&UNIT,
                 SPACE=(1000, (&TAPES)), AVGREC=U
//
//XTRCWK01 DD UNIT=&UNIT,
//
                 SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
//XTRCWK02 DD UNIT=&UNIT.
                 SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
//
//XTRCWK03 DD UNIT=&UNIT,
//
                 SPACE=(350, (&TAPES),, CONTIG, ROUND), AVGREC=U
//XTRIN
           DD DISP=SHR, DSN=&HLQ..GFTAXTR.&RUN
           PEND
//STEP1 EXEC GFTAVMA, HLQ=your.hlq, RUN=XTRACT#1
                                                   <-- OUT HLQ?
//VMA.SYSPRINT DD DISP=(,PASS),DSN=&&VMAPRIN,
// UNIT=SYSDA, SPACE=(CYL, (1,1))
                                                    <-- dsname?
//VMA.XTRIN DD DISP=SHR,dataset=output.from.step.1
//VMA.VMACNTL DD *
JOBNAME (EXCLUDE (SVTS))

↓ Modify VMA Exclusions

REPORT(USAGE)
//*-----
//STEP2 EXEC PGM=IRXJCL.PARM='VTMODEL'
//SYSEXEC DD DSN=SVTS.JCLLIB.NAME.DISP=SHR
                                               <-- VTAPE LIB NAME?
//SYSTSPRT DD SYSOUT=*
             DSN=&&VMAPRIN,DISP=(OLD,DELETE)
//INVMA DD
//* CA-VTAPE ANALYSIS COMMANDS FOLLOW. REMOVE COMMENT * TO USE
//SYSIN DD *
            VSIZE(400)
            OUTPHYVOLSIZE(800) <- 800 indicates 800MB or .8 GB
            VSIZE(800)
            NOREPORTCACHE
            NOREPORTUSAGE
            NOREPORTPHYVOL
            SHOWMIN(0)
            SHOWMAX(99999)
//
```

Example Output Reports

The following sections describe each section of the output from the modeling utility.

Heading and Input Control Statement Listings

The first part of the output contains headings and records the date processing was performed. It also includes a review of supplied input parameters. Messages may also be shown that reflect other options or assumptions made during processing.

```
CA VIRTUAL TAPE MODELING FACILITY ****
      © 1999-2001 Computer Associates ****
      Version 1.2 3/3/1999
Report date: Friday 5 Mar 1999 Time: 15:51:20
______
MDMSG003 MODELER INPUT COMMAND STATEMENTS:
-----8
SHOWMIN(0) SHOWMAX(99999) VSIZE(400) CACHEFACTOR(20)
______
VTAPE MDMSG001 Modeling for CA-Vtape 400 MB virtual volume size -- VSIZE(400) parm
VTAPE MDMSG102 Processing 25 days tape processing input data.
```

Report 1 - Volume Usage Summary Calculations

	- F	REPO	ORT #1	CA-VTAPE	VOLUME USAGE	E SUMMARY CAL	CULATIONS
 Size	MB e Rang	ge	Unique VOLSERS	Total GB	Cumulative VOLSERS	Cumulative GB	Cumulative Avg Vol MB
<=	40	MB	6615	258.4	6615	258.4	40
<=	80	MB	1125	87.9	7740	346.3	46
<=	120	MB	509	59.6	8249	405.9	50
<=	160	MB	1097	171.4	9346	577.3	63
<=	200	MB	1342	262.1	10688	839.5	80
<=	320			437.2	12087	1276.6	108
<=	400			352.7	12990	1629.4	128
<=	520			528.1	14030	2157.5	157
<=	600	MB		427.1	14759	2584.6	179
<=	720		,	800.2	15897	3384.8	218
<=	800			1269.5	17522	4654.3	272
<=	920			729.5	18334	5383.9	301
•	1000			590.8	18939	5974.7	323
	2000			4031.3	21003	10005.9	488
	3000			1579.1	21542	11585.0	551
•	4000			343.8	21630	11928.8	565
•	5000		24	117.2	21654	12046.0	570
>	5000	MB	10	97.6	21664	12095.2	572

This report shows a breakdown of the number of tape volumes in your installation and the number of volumes in megabyte size ranges. It also provides details about how CA-Vtape can be effective in consolidating tape volumes and what sizes of data sets you want CA-Vtape to manage.

The following information is shown:

MB SIZE RANGE

Column indicates the size ranges for each row of analysis calculations. Each row of information shows totals applying to all tapes SMF reported were processed during the logging period analyzed.

UNIQUE VOLSERS Indicates the unique volume serial

numbers the analysis included.
"Unique volsers" refers to the
number of created or newly
processed volumes during the
analysis period. This does not
include tape inventory that was not
created or read during the logging

period.

TOTAL GB The total gigabytes (billions of

bytes) contained on the unique volumes in this size range of tape data. This may be the total bytes to be considered for CA-Vtape cache

or physical volume storage.

CUMULATIVE VOLSERS Contains the total unique volsers in

this size range and all size ranges lower than this range. It is a total of the column up to and including this

row.

CUMULATIVE GB Contains the total unique gigabytes

in this size range and all size ranges lower than this range. It is a total of the column up to and including this

row.

CUMULATIVE MB Indicates the average size of all

unique volumes found in this volume size range. This value is used in the calculation of the cache size needed to hold this physical

volume.

Report 2 - Cache Size Alternatives to Preserve Cache Residence

--REPORT #2 -- CA-VTAPE CACHE SIZE ALTERNATIVES TO PRESERVE CACHE RESIDENCE --

VTAPE MDMMSG016-Increased estimated cache size by a CACHEFACTOR of 20 percent.

I MB	Cumulative	Cumulative	Cumulative	Cumulative	Cumulative	- 1
Size Range	GB 6 hrs	GB 12 hrs	GB 24 hrs	GB 36 hrs	GB 48 hr	i
<= 40 MB	3.1	6.2	12.4	18.6	24.8	- 1
<= 80 MB	i 4.2 i	8.3 i	16.6	i 24.9	33.2	i
<= 120 MB	i 4.9 i	9.7	19.5	i 29.2	39.0	i
<= 160 MB	j 6.9 j	13.9	27.7	41.6	55.4	i
<= 200 MB	10.1	20.1	40.3	60.4	80.6	i
<= 320 MB	15.3	30.6	61.3	91.9	122.6	i
<= 400 MB	19.6	39.1	78.2	117.3	156.4	i
<= 520 MB	j 25.9 j	51.8	103.6	155.3	207.1	i
<= 600 MB	j 31.0 j	62.0	124.1	186.1	248.1	i
<= 720 MB	i 40.6 i	81.2	162.5	243.7	324.9	i
<= 800 MB	55.9	111.7	223.4	335.1	446.8	Ĺ
<= 920 MB	64.6	129.2	258.4	387.6	516.9	Ĺ
<= 1000 MB	71.7	143.4	286.8	430.2	573.6	Ĺ
<= 2000 MB	120.1	240.1	480.3	720.4	960.6	- 1
<= 3000 MB	139.0	278.0	556.1	834.1	1112.2	- 1
<= 4000 MB	143.1	286.3	572.6	858.9	1145.2	Ì
<= 5000 MB	144.6	289.1	578.2	867.3	1156.4	Ì
> 5000 MB	145.1	290.3	580.6	870.9	1161.1	Ì

Shows projections of the size of DASD cache to be set up for different size ranges of CA-Vtape volumes. It helps you to understand the amount of DASD cache needed, depending on how long virtual volumes are to remain on quickly accessible cache storage.

The following information is shown:

MB SIZE RANGE Indicates the size ranges where the row of analysis calculations apply. Each row contains totals for all tapes SMF reported were processed during the logging period analyzed.

RESIDENCE GB x HOURS

(6, 12, 24, 36, 48)

Five columns of estimated cache size requirements are shown. Each row provides cache size details for different tape volume size ranges. Each column shows the cache gigabytes needed to hold that size range in cache for a specified number of hours. The column shows a cumulative total that includes the cache amount required for this size range and the size ranges above it on the report. The longer data sets can be accessed from cache, the faster job throughput. It also consumes less overall overhead and reduces tape recall time. You can use these values to determine how to balance cache size against processing performance.

The value shown in this column can also reflect an additional percentage of cache needed as specified by the CACHEFACTOR(0-100) input parameter. For more information about CACHEFACTOR, see Report 1 - Volume Usage Summary Calculations in this chapter.

Report 3 - Tape Usage Projection and Estimated Cartridge Savings

	- REPORT	#3	CA-VT	APE TA	PE USA	AGE PR	OJECTIO	N AND	ESTIMATED	CARTRI	DGE SAV	INGS	
		GB	Avg	 Curr	Est	Est	Cum	Cum	Cum	Cum	Cum	Pct	Pct
	MB	All	Vol	Phys	Phys	Saved	GB All	Avol	Phys	Est	Saved	Sav-	Total
Si:	ze Range	Vols	MB	Vols	Vols	Vols	Vols	MB	Vols	Vols	Vols	ings	Savngs
<=	40 MB	77	12	6615	43	6572	77	12	6615	43	6572	99.3%	40.2%
<=	80 MB	65	60	1125	37 j	1088	142	71	7740	80	7660	96.7%	46.8%
<=	120 MB	45	92	509	26	483	188	163	8249	105	8144	95.0%	49.8%
<=	160 MB	135	126	1097	76	1021	322	289	9346	181	9165	93.1%	56.0%
<=	200 MB	243	185	1342	136	1206	565	474	10688	317	10371	89.8%	63.4%
<=	320 MB	347	254	1399	195	1204	912 j	728	j 12087 j	512	11575	86.1%	70.8%
<=	400 MB	312	354	903	175	728	1224	1082	12990	687	12303	80.6%	75.2%
<=	520 MB	461	454	1040	259	781	1685	1536	14030	946	13084	75.1%	80.0%
<=	600 MB	394	554	729	221	508	2080	2090	14759	1167	13592	69.6%	83.1%
<=	720 MB	727	654	1138	408	730	2807	2744	15897	1575	14322	64.2%	87.5%
<=	800 MB	1197	754	1625	671	954	4003	3498	17522	2246	15276	58.7%	93.4%
<=	920 MB	677	854	812	380	432	4680	4352	18334	2626	15708	53.2%	96.0%
<= ;	1000MB	564	954	605	316	289	5244	5306	18939	2942	15997	47.7%	97.8%
<= [2000 MB	3032	1504	2064	1701 j	363	8276	6810	j 21003 j	4643	16360	17.6%	100.0%
<=]	3000 MB	1318	2504	539	539	0	9594	9314	21542	5182	16360	0.0%	100.0%
<= .	4000 MB	301	3504	88	88	0	9895	12818	21630	5270	16360	0.0%	100.0%
<=	5000 MB	106	4504	24 [24	0	10000	17322	21654	5294	16360	0.0%	100.0%
>	5000 MB	49 أ	5005 i	10 i	10	01	10049 i	22326	1 21664 1	5304	16360	0.0%	100.0%

Projects the potential tape cartridge savings achieved when CA-Vtape is used to manage different size ranges of data sets. It shows the dramatic savings achieved by using CA-Vtape.

The following information is shown:

MB SIZE RANGE	Indicates size ranges the row of analysis calculations applies to. Each row of information contains totals applying to all tapes SMF reported were processed during the logging period analyzed.
GB ALL VOLS	Shows the total gigabytes on all tape volumes in this size range. This is the uncompressed size of data stored in

cache and on physical backstor volumes.

AVG VOL MB Indicates the average size of all unique

volumes found in this volume size range. This value is used in the calculation of cache size needed to hold physical

volume data.

CURRENT PHYS Contain

VOLS

Contains the original number of volumes analyzed by SMF for all volumes in this size range. If only physical volumes were included in the analysis, this is a total of physical volumes used without CA-

Vtape.

EST PHYS VOLS Based on the specified uncompressed

output physical volume size, this is the number of physical CA-Vtape backstor volumes that are estimated to hold all stacked virtual volumes managed.

EST SAVED VOLS Shows the total reduction in physical tape volumes once they are stacked by CA-Vtape. This reflects a savings in tape storage slots, cartridges, floorspace, and

tape management costs.

CUM GB ALL

VOLS

Shows a cumulative column total of the earlier column in the report labeled GB

ALLVOLS.

CUM AVOL MB Shows a cumulative column total of the

earlier column in the report labeled

AVERAGE VOL MB.

CUM PHYS

VOLS

Shows a cumulative column total of the earlier column in the report labeled

CURRENT PHYSICAL VOLS.

CUM EST VOLS Shows a cumulative column total of the

earlier column in the report labeled ESTIMATED PHYSICAL VOLS.

CUM SAVED

VOLS

Shows a cumulative column total of the

earlier column in the report labeled

ESTIMATED SAVED VOLS.

PCT SAVINGS

Shows the percentage of savings in volumes for the size range for this report line.

PCT TOTAL SAVINGS

Shows the percentage of the total savings in physical volumes for the size range for the current line and all above lines. This percentage differs from the PCT SAVINGS column in that it is a percentage of savings that this size range participates in the total CA-Vtape projected savings.

For example, a value of 20% in this column indicates that this size range and above ranges comprise 20% of total projected savings. Use this column to help find at what size range the savings of physical volumes becomes less substantial.

Chapter

Customer Support

Software that Manages eBusiness

Computer Associates is the company that provides software that manages eBusiness and we see storage management as a crucial part of managing eBusiness. For more than 20 years, Computer Associates has led the market with data storage management tools. Today, our products can help ensure that your missioncritical storage systems, including the revenue producing data they capture and the applications that work with that data are continuously available.

Customer Support

CA-Vtape provides extensive documentation. In addition to the Getting Started, the documentation set includes the User Guide and the Message Guide. These manuals describe how to run the CA-Vtape application. All 3 manuals are delivered on the product CD in Adobe's Portable Document Format (PDF).

Getting Technical Support for CA-Vtape

For further technical assistance with this product, please contact Computer Associates Technical Support using our web-based, self-help support system at http://esupport.ca.com. Technical support is available 24 hours a day, seven days a week.

You can also initiate a support incident on the Computer Associates support system, and track existing incidents on this site as well, by following links to StarTCC.

Note: You must register for the StarTCC site and be a licensed CA-Vtape user in good standing to use it. You need your site ID to register.

Computer Associates also supports an Open Forum for client discussions about CA-Vtape related issues. The Open Forum is modeled on a News Group. Register for the CA-Vtape Forum from the eSupport site http://esupport.ca.com/ by following the links to Open Forums.

To view the documentation of CA-Vtape, see http://support.ca.com/public/sams/samssupp.html. It contains books and release notes that you can view online or download to your PC.

Computer Associates customers also benefit from:

- Local user group meetings that provide customers with upto-date information on product enhancements.
- The annual CA-World customer conference is packed with information, training and interaction with other customers and our product developers.
- To put the power of Computer Associates to work for you, contact your nearest Computer Associates office or visit our Web site today at http://ca.com/.

Index

allocation cache new scratch, 1-10 virtual management, 1-10 analysis VMA and CA-Vtape, 3-7 analyze how much data to, 3-4 analyzer, IBM's volume mount, 3-3 architecture, CA-Vtape, 1-7 automated virtual volume stacking, 1-16 В back store requirements, 2-1 buffer disk pool status, 1-14 pool and stacking management, DASD, 1-17 copies pool, disk, 1-9

C

```
size alternatives, 3-2
    size, report example, 3-15
CACHEFACTOR(0-100), 3-10
calculations, volume usage summary, 3-1
cartridge
    estimated savings, 3-2
catalog, processing ICF, 1-14
CA-Vtape
    technical support, 4-2
condition
    end of volume, 1-12
    other processing, 1-15
    special processing, 1-15
considerations, installation, 2-1
consolidation of SMF data, VMA, 3-5
control statement listings, 3-12
    duplex, 1-18
    export, 1-18
customer support, 4-1
```

D DASD buffer pool and stacking IBM's management, 1-17 volume mount analyzer, 3-3 data ICF catalog processing, 1-14 to analyze, how much, 3-4 input VMA consolidation of SMF, 3-5 parameters, 3-8 data sets, 1-8 record types, SMF, 3-4 retrieving existing, 1-11 install data space, 1-8 considerations, 2-1 processing, 1-13 planning for, 3-1 requirements, 2-1 devices, virtual tape, 1-7 disaster recovery, 1-18 disk buffer pool, 1-9 pool status, 1-14 JCL setup duplex copies, 1-18 VMAJCL1, 3-5 VMAJCL2, 3-7 Е end of volume, condition, 1-12 listings, 3-12 existing data sets, retrieving, 1-11 export copies, 1-18 M G modeling, running CA-Vtape, 3-4 mount analyzer, IBM's volume, 3-3 groups, stacking, 1-17

multi-file

processing to virtual volumes, 1-12

multi-volume output processing, 1-12

hardware requirements, 2-1 how CA-Vtape works, 1-3

Н

N	multi-file to virtual volumes, 1-12 multi-volume output, 1-12			
new (scratch) allocations, 1-10	_			
NOREPORTCACHE, 3-9	R			
NOREPORTPHYVOL, 3-9				
NOREPORTUSAGE, 3-9	readability outside CA-Vtape, 1-19			
0	record types SMF input, 3-4			
	recovery, disaster, 1-18			
output processing multi-volume, 1-12 report examples, 3-12	recycle report, 1-20 simulation, 1-20			
outside CA-Vtape, readability, 1-19 overview CA-Vtape, 1-2	recycling process, 1-19 tape, 1-19			
P	release 2.0 documentation changes, 1-24 new features, 1-21 What's new, 1-21			
parameters input, 3-8 pool disk buffer, 1-9 stacking management, DASD buffer, 1-	report cache residence example, 3-15 cartridge savings, example, 3-17 example output, 3-12 recycle, 1-20 tape usage projection example, 3-17 volume usage summary, 3-13			
17 statusdisk buffer, 1-14 preserve cache residence, 3-2 report example, 3-15	reporting simulation, recycle, 1-20 VMA and CA-Vtape analysis, 3-7			
process recycling, 1-19 stacking, 1-16 processing	requirements back store, 2-1 hardware, 2-1 install, 2-1 system, 2-1			
close tape mark, 1-12 conditions, other, 1-15 conditions, special, 1-15	residence cache, 3-2			
data space, 1-13 ICF catalog, 1-14	retrieving existing data sets, 1-11			

running	<u>T</u>				
CA-Vtape modeling, 3-4					
savings estimated cartridge, 3-2 tape usage projection, 3-2	tape mark, processing, 1-12 recycle, 1-19 usage efficiency, 1-6 usage projection, 3-2 virtual devices, 1-7 TCO, lower, 1-6				
scratch new allocations, 1-10	technical support for CA-Vtape, 4-2				
SHOWINLINES, 3-9					
SHOWMAX(99999), 3-10	V				
SHOWMIN(0), 3-9					
simulation and reporting, recycle, 1-20 SMF data, 3-5 input record types, 3-4 special processing conditions, 1-15 specifications, high level, 1-7 stacking automated virtual volume, 1-16 groups, 1-17 management, 1-17 management, DASD buffer pool, 1-17 process, 1-16 statement listings, 3-12 status, disk buffer pool, 1-14	virtual allocation management, 1-10 tape devices, 1-7 volume stacking, automated, 1-16 volumes, 1-7 volumes, multi-file processing to, 1-12 VMA analysis and reporting, 3-7 consolidation of SMF data, 3-5 volume end of, 1-12 mount analyzer, IBM's, 3-3 multi-file processing to virtual, 1-12 output processing multi-, 1-12 stacking, automated virtual, 1-16 usage report example, 3-13 virtual, 1-7				
support customer, 4-1 technical, for CA-Vtape, 4-2	W				
•	**				
system requirements, 2-1	What's new in release 2.0, 1-21				

